Calibration of the Far Infrared Absolute Spectrophotometer (FIRAS) on the Cosmic Background Explorer (*COBE*[1])

D. J. Fixsen[2], E. S. Cheng[3], D. A. Cottingham[4], R. E. Eplee, Jr.[5], T. Hewagama[6], R. B. Isaacman[5], K. A. Jensen[6], J. C. Mather[3], D. L. Massa[2], S. S. Meyer[7], P. D. Noerdlinger[2], S. M. Read[6], L. P. Rosen[6], R. A. Shafer[3], A. R. Trenholme[5], R. Weiss[7], C. L. Bennett[3], N. W. Boggess[3], D. T. Wilkinson[8], and E. L. Wright[9]

ABSTRACT

The FIRAS instrument on the *COBE* satellite was designed to accurately measure the spectrum of the cosmic microwave background radiation (CMBR) in the frequency range 1 to 95 cm^{-1} with an angular resolution of 7°. We describe the calibration of this instrument, including the method of obtaining calibration data, reduction of data, the instrument model, fitting the model to the calibration data, and application of the resulting model solution to sky observations. The instrument model fits well for calibration data which resemble sky conditions. The method of propagating detector noise through the calibration process to yield a covariance matrix of the calibrated sky data is described. The final uncertainties are variable both in frequency and position, but for a typical calibrated sky 2.6° square pixel and 0.7 cm^{-1} spectral element the random detector noise limit is of order of a few times

---

[1]The National Aeronautics and Space Administration/Goddard Space Flight Center (NASA/GSFC) is responsible for the design, development, and operation of the Cosmic Background Explorer (*COBE*). Scientific guidance is provided by the *COBE* Science Working Group. GSFC is also responsible for the development of the analysis software and for the production of the mission data sets.

[2]Applied Research Corporation, Code 685.3, NASA/GSFC, Greenbelt MD 20771

[3]NASA Goddard Space Flight Center, Code 685, Greenbelt MD 20771

[4]Universities Space Research Association, Code 685.3, NASA/GSFC, Greenbelt MD 20771

[5]General Sciences Corporation, Code 685.3, NASA/GSFC, Greenbelt MD 20771

[6]Hughes-STX, Code 685.9, NASA/GSFC, Greenbelt MD 20771

[7]MIT Dept. of Physics, Room 20F-001, Cambridge MA 02139

[8]Joseph Henry Laboratories of Physics, Princeton Univ., Princeton NJ 08540

[9]UCLA Astronomy Dept., Los Angeles CA 90024-1562

10^{-7}erg cm^{-2} s^{-1} sr^{-1} cm for 2 to 20 cm^{-1}, and the difference between the sky and the best fit cosmic blackbody can be measured with a gain uncertainty of less than 3%.

Subject headings: cosmology: cosmic microwave background — instrumentation:interferometers

1. Introduction

The FIRAS (Far Infrared Absolute Spectrophotometer) instrument aboard the *COBE* (Cosmic Background Explorer) satellite (Mather *et al.* 1990, Boggess *et al.* 1992), launched November 18, 1989, was designed to measure the cosmic microwave background radiation (CMBR) spectrum. The instrument was designed for higher accuracy and sensitivity than previously obtained in the frequency range 1 to 95 cm^{-1}, with a goal of 0.1% of the peak brightness of the CMBR for each spectral element of 6144 equal area pixels on the sky. New FIRAS results show that the CMBR has a blackbody form within 0.03% of the peak brightness, with a temperature of 2.726 ± 0.010 K (95% CL systematic) (Mather *et al.* 1994), with a negligible statistical error of only 0.00001 K (95% CL), and that the spectrum of the dipole anisotropy of the CMBR is a Doppler-shifted blackbody at the same temperature as the CMBR to within 0.02% of the peak brightness (Fixsen *et al.* 1994). Atomic and ionic fine structure lines of [C I], [C II], and [N II], molecular rotational transitions of CO, and dust emissions from the interstellar medium have also been mapped throughout our Galaxy (Wright *et al.* 1991, Bennett and Hinshaw 1993).

The FIRAS shares a liquid helium cryostat with the Diffuse Infrared Background Experiment (DIRBE), but no electrical or optical crosstalk between them has been detected. The cryostat contained ~ 600 liters of liquid helium at launch, and provided a thermal reservoir at 1.4 K for 10 months of operation. The FIRAS line of sight is along the symmetry and spin axis of the spacecraft, and is maintained away from the Earth and approximately 94° from the Sun during the mission. The orbit is a 99° inclination, 900 km altitude, sun-synchronous orbit with a 6 PM ascending node, so that the orbit plane is always approximately perpendicular to the direction of the Sun. The precession of the orbit plane, caused by the equatorial bulge of the Earth, is used to scan the instrument line of sight continuously around the sky. During the 10 month helium lifetime, the FIRAS line of sight scanned the sky 1.7 times. The spacecraft spins about its symmetry axis but this does not affect the FIRAS pointing.

In this paper we describe the method of calibration of the FIRAS, including the mathematical model of the FIRAS emissions, the detector model, and the data reduction process. This calibration has reached a reasonable degree of self-consistency after the recognition and modeling of certain imperfections of the instrument. These imperfections include a small temperature calibration drift in the internal reference body, optical harmonics due to multipass phenomena, and sidebands in the spectra induced by vibrations of the movable mirror.

2. Instrument

2.1. Interferometer

The FIRAS is a symmetric polarizing Michelson Fourier transform spectrophotometer with two input and two output ports (see Figure 1). Using bolometric detectors (Mather 1984a), it is sensitive to radiation from 1 to 95 cm^{-1}. The difference between the spectra of the radiation coming into the two input horns is measured by detectors at each of the two outputs (arbitrarily labeled "right" and "left"). Each input beam is split by a wire grid polarizer such that the horizontal polarization is transmitted and the vertical polarization is reflected. After being reflected from the collimating mirrors the beams are again split by another polarizer at 45° (about the axis of the beam) relative to the first polarizer, then sent down the two arms of an interferometer. After reflection from dihedral mirrors, the beams return to the splitting element with a polarization rotated by 90°, where they are recombined. They reflect again from the collimating mirrors, and return to the input polarizer where they split into the two outputs. The division of light between the two outputs is governed by the difference in path length between the two arms of the interferometer and the wavelength of the incoming radiation.

The instrument has a nominal étendue, or light gathering power (cosine-weighted area-solid angle product $A\Omega = \int dA \cos\theta d\Omega$), of $A\Omega = 1.5$ cm^2 sr, where A is the collecting area, and Ω is the solid angle received. This is defined physically by the most restrictive aperture in the system, the throat of the input horn antenna, and is $A\Omega = \pi^2 r^2$, where r is the radius of the aperture. The number of independent spatial modes of the radiation field received is $N_s = A\Omega/\lambda^2$ and each such mode can have two polarizations. Multiple refocusing mirrors are needed to keep the beam from diverging beyond the size of the optical elements. The beam divergence sets a limit of $\sim 0.8\%$ on the spectral resolution.

The dihedral reflectors for both arms of the interferometer are mounted on the two ends of a single mirror transport mechanism (MTM). The MTM scans slowly in the forward direction, with a choice of two speeds, and then flies back to the starting point. Data taking commences when the MTM passes a reference mark on a scale, after most of the mechanical transients excited by the abrupt velocity change have decayed. No data are taken in the reverse direction. The MTM has flexible leaf spring bearings, so the motion is smooth. The MTM stroke length determines the spectral resolution of the instrument, with the best unapodized spectral resolution given by $\Delta\nu = 1/(2\delta_{max})$, where δ_{max} is the maximum path difference between the two arms. The maximum path difference is 5.85 cm and the best spectral resolution is $\Delta\nu = 0.0855$ cm^{-1}, based on the assumption of symmetry of the interferogram about the zero path difference point. This limits the resolution at low frequencies, but the 0.8% beam divergence is the limiting factor at the highest frequencies (near 95 cm^{-1}).

A dichroic filter separates each output into a "low" frequency (1 to 20 cm^{-1}) and a "high" frequency (20 to 95 cm^{-1}) channel at each output port, making a total of four channels ("right high", "left low", etc.). (The "high" frequency channels have some sensitivity to radiation in the "low" band, but the "low" channels have negligible sensitivity to radiation in the "high" band.) Each channel has a bolometer detector which, depending on the length and speed of the MTM stroke, was sampled from 1024 to 6144 times per stroke. The data were digitally filtered and averaged on board the spacecraft so 512 points were telemetered to the ground for each interferogram regardless of scan length or speed. There were two scan lengths and two scan speeds, making possible a total of 16 sets of data among the four detectors.

In this paper we will discuss only the short slow and long fast scan modes, since $\sim$ 88% of the data were taken in these modes. There are $\sim$234,000 long fast interferograms and $\sim$275,000 short slow interferograms for each detector, each of which is the onboard average of either four (long scan) or 16 (short scan) MTM strokes. Collection times for these averaged interferograms were 53 sec for the short slow mode and 31 sec for the long fast mode. Figure 2 shows a typical interferogram.

All of the calibration measurements were made assuming that each emitter in the field of view (*e.g.* the sky horn) is at a well defined and known temperature. To make the instrument fulfill this assumption, each component was made with a low internal thermal resistance and a comparatively high thermal resistance to other parts and to the dewar. This was particularly important in the case of the external calibrator (discussed in the next section), the ultimate calibration standard for the FIRAS.

2.2. Calibrators

The FIRAS was calibrated on the ground and in flight by a movable external blackbody calibrator, hereafter called the Xcal. When in place, the Xcal fills the entire aperture of the sky horn. The other input horn, called the reference horn, is filled by a fixed black body calibrator, called the internal calibrator or Ical. Since the instrument is differential, the output of the instrument is zero when the emitted spectrum from the Ical matches that of the sky (or Xcal). This null condition greatly reduces the sensitivity of the experiment to a host of potential systematic errors.

The Xcal (see Figure 3) resembles a trumpet mute with a 25° groove angle, with a front surface of iron loaded epoxy (Emerson and Cumming CR-110; see Hemmati, Mather, and Eichhorn 1985) designed to be absorptive at microwave frequencies. The support for the absorbing material is copper and copper foil, to provide high thermal conductivity between different parts of the Xcal. The heater for the Xcal was mounted at the joint between this copper backing and the thermal isolator which supported the Xcal, so no heat flow passes through the body of the Xcal in steady state. As a result, calculations show that the temperature inhomogeneity in the Xcal cannot exceed 1 mK at 2.7 K. The support arm is also used to move the Xcal into and out of the sky horn.

When the Xcal is in the sky horn it does not quite touch it. There is a 0.6 mm gap between the edge of the Xcal and the horn, so that the Xcal and the sky horn can be at different temperatures. Although the gap is near the flare of the horn and not in the direct line of sight from the detectors, it would result in undesirable leakage at long wavelengths because of diffraction. To ensure a good optical seal at all wavelengths, two ranks of aluminized Kapton leaves attached to the Xcal make a flexible contact with the horn.

Several measurements were made to test how closely the Xcal approximates an ideal blackbody. The reflectance of the Xcal in the horn was measured at several frequencies near 1 and 3 cm^{-1} and the largest reflections were less than 3×10^{-5}. An upper limit for leakage around the Xcal was determined in ground tests with a warm cryostat dome by comparing signals with the Xcal in and out of the horn. Leakage is less than 1.5×10^{-4} in the range $5 < \nu < 20$ cm^{-1} and 6.0×10^{-5} in the range $25 < \nu < 50$ cm^{-1}. The transmission of the material is significantly higher at cryogenic temperatures than at room temperature (Hemmati, Mather, and Eichhorn 1985). Calculations showed that transmission through the material might be important at the very longest wavelengths, so the back of the Xcal was covered with opaque foil. The Kapton leaves sealing the gap between sky horn and Xcal were tested by gradually withdrawing the Xcal from the horn. No effect could be seen in flight until it had moved 1.2 cm. In short, the Xcal is an excellent approximation to an ideal blackbody that fully fills the field of view.

The Ical construction is somewhat different. It is permanently mounted at the mouth of the reference horn and optically sealed by a labyrinth of absorbing material. The Ical could have (and perhaps should have) also been a good blackbody. However, it was thought that some nulling of the emission from dust might be required, so the design included two "hot spots" that could be heated independently of the main reference body. The Ical has a metal cup (∼1 cm) on its face containing "hot spot" heaters which could be raised to a much higher temperature. The result is that the Ical has a $\sim 4\%$ reflectance with small interference ripples as a function of frequency. This departure of the Ical from a true blackbody is one of the elements of the calibration model.

2.3. Horn Inputs

Both machined aluminum horns are connected to the other parts of the instrument by low thermal conductivity fiberglass supports. Each of these also has a heater near the point of attachment, so no steady state heat flow exists in the body of either horn.

The sky horn of the FIRAS provides a 7° field of view, with low sidelobes (Mather, Toral and Hemmati 1986). It is a non-imaging parabolic concentrator with a curved flare at the mouth to attenuate sidelobes from far off axis. Measurement of regions of the sky near the Moon partially confirm the sidelobe pattern measured before launch, and are less than $-38\,\mathrm{dB}$ beyond 15° from the center of the beam. Preliminary results are shown in Figure 4, along with preflight measurements at 1 and 1.77 cm^{-1}. The flight antenna pattern continues to be studied by analysing data taken at various angles from the Moon. The Moon is clearly detectable at angles out to 15° from the line of sight.

Ideally the two input horns should have zero emission and be symmetric. In fact the emission is $\sim 0.012 + 0.0015 \times (\nu/1\,\mathrm{cm}^{-1})$ of the emission from the Xcal and is symmetric at low frequencies. Spacecraft limitations forced the reference horn to be much smaller than the sky horn. The same effective length-to-diameter ratio matched them well at low frequencies but not so well at high frequencies. Since they are kept at 2.7 K for sky observations, the high frequency difference is not critical. Calculations including the anomalous skin effect show qualitatively the same features. However, the measured emission is higher than predicted, particularly at the lowest frequencies. The horns do not have the same pattern of emission as a function of position and angle within the beam as do the blackbodies, and so do not necessarily produce identical optical responses.

2.4. Thermal Design

The parts of FIRAS that are in the beam were kept cold and at stable temperatures, so that there would be as little internally generated signal as possible. The primary heat sources are the MTM, the detector JFET amplifiers, the control heaters for the calibrators and horn antennas, and sources within the DIRBE instrument. Ultimately all the heat reaches the helium cryostat and is carried away by helium vapor.

Several germanium resistance thermometers (GRTs) were attached to the Xcal, the Ical and the two horns (collectively called the "controllables"), and one GRT was used in a feedback loop with each heater to actively control the respective temperatures. The control circuit was much faster than the thermal time constant of the controllables (minutes), and no oscillations were seen. All thermometers were calibrated at Goddard Space Flight Center relative to a NIST calibrated standard GRT, nominally to 1 mK accuracy.

Other parts, made mainly of aluminum, also have thermometers, but not heaters. These include the MTM, the bolometer housings, and three parabolic collimating mirrors. In normal operation the temperature for all of these parts was stable and between 1.5 K and 2.7 K.

The temperatures of the four principal instrumental sources of radiation, the external and internal calibrators, and the sky and reference horns, could all be controlled over the range 2 to 25 K. While under control, these temperatures were generally stable to within the monitoring thermometer resolution, about $0.2\,\mathrm{mK}\,(T/2.7\,\mathrm{K})^3$. Temperature changes were caused by commanded temperature changes, changes of the scan length of the MTM, motion of the Xcal, passage through the aurora, illumination by the Earth, running the MTM into its stop, and calibration or annealing operations of the DIRBE instrument. During calibration, the Xcal was moved into the sky horn, filling the sky aperture, and all four controllable sources were commanded to a variety of temperature combinations. These data are used to derive a calibration model, as described below. In contrast, when observing the sky (*i.e.* when the Xcal was out of the field of view), the Ical and the two horn temperatures were commanded to near 2.7 K to null the signal on the detectors and to remove the dependence on the emissivity of the input horns. Tests were made to determine whether the Xcal temperature made any difference when it was outside the field of view. The correlation of the Xcal temperature with the data implies that the contamination of the data is less than $\Delta I_\nu = 4 \times 10^{-9}$ erg cm^{-2} s^{-1} sr^{-1} cm over the full mission.

There were two important problems. One was that the thermometers on both the Ical and the Xcal did not all agree. In fact, the disagreement among different Xcal thermometers was 3 mK at 2.7 K. An explanation for this discrepancy is that there were shifts in the

GRT calibration between the initial ground calibration and flight. This explanation is supported by remeasurement of spare GRTs from the same batches used in the flight instruments, and by measurement of gradients in an identical spare Xcal.

The disagreement between the two Ical thermometers was 18 mK at 2.7 K. The heat sinking of the Ical thermometer leads was inadequate, and some of the applied heat flowed through part of the Ical. As a result there is a measured gradient in the device that depends on external conditions, and tests have been performed to quantify this effect.

Differences in the horn thermometers showed transients only when the temperatures were changed, and these could be modeled well using the thermal conductance and heat capacity of aluminum. Furthermore, emission from the horns and other parts of the instrument was much smaller than that from the calibrators so the measurement is less critical.

The second problem was that transient events heated the MTM and other parts of the instrument by $\sim$ 1 K. It was some time into the flight before the problem was deciphered and ameliorating commands were sent. At times, charged particles affected the MTM zero reference mark photodetector so that the MTM was driven past its electronic stop and into a physical stop. This did not damage the MTM, but the linear motor driving it dissipated a power of $\sim$200 mW for a few tens of seconds until a timer circuit restarted the scans. This heat then diffused into other parts of the instrument, at times heating them to temperatures as high as 5 K. Most occurrences were in the high radiation region known as the South Atlantic Anomaly. After Dec. 9, 1989 the MTM was held in position during South Atlantic Anomaly passages, but this still increased the heat dissipation of the MTM motor, so after Mar. 21, 1990 the MTM was turned off during those passages. Because of these events in the early part of the flight the temperatures of the mirrors and wire grids cannot be treated as constant. Thus radiation from the MTM and mirrors are included in the calibration model. These MTM events provided the data to deduce the effects of the mirrors and grids and also provided one of the severest tests of the calibration model.

2.5. Electronics and Microprocessor

Some signal processing was done on the spacecraft, to reduce the data transmission rate to the ground. First, AC coupling was used to reduce the scale of the bolometer signals and permit high preamplifier gains. To first order the bolometer acts as a low pass filter (Mather 1982, 1984a, 1984b), so an electronic high frequency boost was used to compensate for it. Additional low pass filtering was used to eliminate high frequency noise, and an adjustable gain was used to amplify the signal to an appropriate level. The resulting electronics transfer function was a five pole Bessel lowpass filter with a $-3\,\mathrm{dB}$ bandwidth of 100 Hz and a six pole highpass filter that is $-3\,\mathrm{dB}$ at 0.3 Hz. The signal was then digitized at a mean frequency of 681.43 Hz, using pulses from a phase lock loop driven by marks on the glass reference scale on the mirror mechanism. Short slow scans were sampled 1536 times per scan, and long fast scans 4096 times per scan.

A separate microprocessor for each detector was used to reduce each interferogram to 512 points, using a digital low pass filter and then averaging the output points in groups of 2, 3, 8, or 12 as appropriate to the scan speed and length and the frequency range of the detector. The number of data points averaged in each scan mode/frequency combination is summarized in Table 1. For the high frequency channels the number of points was 2 for fast and 3 for slow scans, so only 1/4 of each long scan could be recorded. For low frequency channels, the short scans were treated the same as the high frequency channels, but averages of 8 and 12 points were used for long scans to record the entire length of the stroke. The net result of all of this filtering is shown in Figure 5. This transfer function was verified on the ground by direct measurement of the electronics.

As the digital filtering was done, a glitch recognition algorithm was applied to the data as well, and samples with such glitches were marked for use in the ground data processing. These glitches arise from the impact of cosmic rays on the detectors, and are most easily detected when the full data stream is available. The mean rate of glitches found was $\sim 2\,\mathrm{s}^{-1}$ outside the South Atlantic Anomaly.

The microprocessors also averaged 4 long or 16 short successive scans together in groups before transmission, and buffered the data for telemetry. Each number received on the ground is therefore already an average of many digitizations (8 to 48). The interferometer scans were made asynchronous from the telemetry clock to avoid coherent interference from the telemetry system to the detectors.

Other data about the FIRAS and *COBE* were also sent to the ground, such as temperatures of various components, voltages, currents, verification of commands, and the attitude of the spacecraft. This information was later used to reconstruct the state of the FIRAS and to determine its line of sight. Associating the engineering data with the asynchronous mirror scans required considerable care.

If both the sky and the Ical were blackbodies, and the interferometer were perfectly symmetrical, one could in principle null the signal from the former simply by adjusting the temperature of the latter. The temperature of the CMBR could then be read from the reference body thermometers. Unfortunately, neither of those conditions prevails. Our Ical and instrument asymmetry combine to produce a net reflectance of $\sim 4\%$, and Galactic emission from gas and dust contributes to the observed signal. To measure these effects, we must calibrate the instrument.

3. Calibration

3.1. Calibration Overview

Our measured signal (i.e. our interferogram amplitude as a function of path difference) is the convolution of the electronics transfer function with the bolometer output voltage. The bolometer output in turn is a product of the bolometer responsivity and the incident power. The power depends on the spectrum S_k of the kth radiation source and the product a_k of the étendue and the emissivity associated with that source. These functions, determined from the calibration data, are used to recover the sky spectrum.

The detector is nonlinear for large signals, but in the regime of interest a linear approximation can be used in our calibration models. This allows us to perform the entire analysis with linear operations in the frequency domain. Our spectrum Y_f is the inverse Fourier transform of the interferogram (I), giving

$$Y_f = \int_{-\infty}^{+\infty} I(\delta) e^{-2\pi i \nu \delta} d\delta = Z(f) B(f) \sum_k a_k(\nu) S_k(\nu), \qquad (1)$$

where Z is the electronics transfer function, B is the complex bolometer responsivity, and ν is the optical frequency in cm^{-1}, related to the data frequency by $f = v\nu$, where v is the effective rate of change of optical path difference caused by the MTM movement. Since this rate is constant for a given scan mode (0.7827 cm s^{-1} for slow and 1.174 cm s^{-1} for fast), there is a one to one correspondence between ν and f. Hereafter we will use f to denote either the optical or data frequency and depend on the context to distinguish the cases. For compactness of notation, subscripts are used to write $1/Z(f) = Z_f^{-1}$, $1/B(f) = B_f^{-1}$, $a_k(\nu) = a_{fk}$.

The functions a_{fk} have the same dimensions as the étendue and are thus defined here as effective étendues. With the normalization chosen here, the perfect Michelson interferometer would have 1/2 of the design étendue of 1.5 cm^2 sr for the Xcal or sky. Graphs of the transfer functions for one set of calibration data are shown in Figure 6. Known optical losses account for most of the inefficiencies of the system.

Emitting objects other than the sky, i.e. parts of the instrument, are approximated by Planck functions since they are nearly isothermal (sections 2.2-2.4). Hence, using the notations described above, and P for the Planck function $P_f(T) = 2hc^2\nu^3/(e^{h\nu c/kT} - 1)$, solving equation (1) for the sky spectrum gives

$$S_{fsky} = e^{if\psi} a_{fsky}^{-1} \left[\frac{Y_f}{B_f Z_f} - \sum_k a_{fk} P_f(T_k) - \text{corrections} \right] . \tag{2}$$

The application of the model to sky data is briefly summarized in the following five steps.

(1) A model of the analog electronics and digital processing filters, Z_f, is used to remove the transfer function of the electronics.

(2) The state of the bolometer is found from the temperatures and other engineering data. Then a linear model of the bolometer, B_f, (Mather 1984b) is used to infer the radiation spectrum as seen at the detector.

(3) Emissions from all of the known internal sources in the beam are subtracted from the data, including corrections for optically induced errors such as multipass phenomena and vibration-induced sampling errors.

(4) The result is divided by the effective étendue $a_{fsky} \equiv a_{fx}$, where x denotes Xcal.

(5) Finally, a small phase adjustment ψ (section 3.5) is made to produce the net result S_{fsky}, which is presumed to be the sky spectrum associated with the point of interest.

It is important to remember that to first order, at frequencies characteristic of the CMBR, this is a null experiment: we are looking at $\sim 1\%$ differences from a blackbody spectrum at 2.73 K. Thus, a 1% error in a_{fsky} produces a relative error of only 10^{-4} in the measurement of the radiation. Both the reflectivity of the Xcal and our detector noise allow us in principle to achieve this absolute photometric accuracy, and so our goal is to achieve 1% absolute calibration accuracy. For high frequency data (>20 cm^{-1}), however, the picture is different. Here, there is no nulling signal from the reference side, so for these data the calibration errors propagate directly to the measurement and dominate the final errors.

3.2. Fourier Transforms and Apodization

While equation (1) is general and continuous, a numerical evaluation is an approximation of the Fourier transform $\mathcal{F}^{-1}(I)$ requiring a finite set of samples over a finite range of path difference δ. The finite range of path differences is represented by an apodization function $A(\delta)$ which multiplies the interferogram and must be zero outside the measured range. We have chosen not to center the interferogram peak in the measured range, in order to obtain greater spectral resolution. The inverse Fourier transform of the apodization function gives the spectral resolution window function. The apodization functions used for this work are illustrated in Figure 7, along with the real and imaginary parts of the spectral window function. The chosen function is the product of two factors, $A_1(\delta)A_2(\delta)$. The first is symmetrical about $\delta = 0$ and declines smoothly to zero at the long end of the interferogram. The second declines smoothly to zero at the short end of the interferogram, has $dA_2(\delta)/d\delta = 0$ at $\delta = 0$, and has $A_2(\delta) + A_2(-\delta) = 2$. This factor A_2 enables the use of both sides of the interferogram peak for broad band photometry and phase determination, while eliminating discontinuities associated with the truncation of the interferogram at the short end.

The interferograms have random detector noise added to the true infrared signals. If the noise comes from a stationary process, then the inverse Fourier transform of the time series will be uncorrelated. In other words, the noise at each different frequency will be an independent Gaussian random variable, with independent real and imaginary parts.

$$\langle N_j N_k^* \rangle = \sum_{lm} e^{-2\pi i(jl-km)/N} \langle I_l I_m \rangle = \sigma_j^2 \delta_{jk}, \tag{3}$$

where N_j is the noise on spectrum point index j and I_l is the noise on interferogram point l. Direct checks on the data indicate that $\langle N_j N_k^* \rangle$, where $k \neq j$, have less than 1% of the total power.

However, apodization of the interferograms prior to inverse Fourier transformation produces a correlation of the noises at different frequencies. The apodized correlation function in the frequency domain is then:

$$\langle N_j N_k^* \rangle = \sum_{lm} e^{2\pi i(jl-km)/N} A_l A_m \langle I_l I_m \rangle \simeq \sigma_j \sigma_k \sum_{lm} e^{2\pi i(jl-km)/N} A_l A_m, \tag{4}$$

For the purpose of deriving the calibration parameters, we ignore even these correlations between frequencies, as they are only $\sim 5\%$ of the diagonal values and the systematic errors are larger.

Since the interferogram is real, its complex inverse Fourier transform is Hermitian; thus half of the resulting spectrum is redundant and is not calculated. We evaluate the finite Fourier transform only at the frequencies $\nu = m/L$, where L is the total MTM stroke length, with m from 1 to 256. After phase correction, the spectra are expected to be real except for narrow spectral features. That is, most of the antisymmetrical parts of the interferogram are systematic features which have been removed by the calibration.

3.3. Optical model

The optical model is simple and linear. The modulated power absorbed by the detector is $R_f = \sum_k a_{fk} P_f(T_k)$, where k runs over the list of seven components: Xcal (when in the sky horn), Ical, sky horn, reference horn, dihedral mirror, collimator mirror and grids, and bolometer. When the Xcal is out of the sky horn, the sky (with an unknown spectrum S_f) replaces the $P_f(T_x)$. Note that since the instrument is differential, the a_k for the Ical and reference horn are negative relative to the inputs on the other side. Of course, the components do not have negative étendues but their signals are modulated 180° out of phase by the motion of the MTM.

3.4. Emitters

The mirrors, beam splitters, and concentrators that define the optics also emit infrared radiation, some of which is modulated and detected by the instrument. Their temperatures were not regulated like the horns and calibrators, but representative locations were monitored. These include the dihedral mirrors on the MTM, the collimating mirrors and wire grids, and all four detector housings. If their temperatures were always constant it would not be possible to separate their effects, but there were temperature fluctuations. Since the instrument was not isothermal the measured temperatures are not necessarily the actual temperatures of the important emitters. Our approach is to model radiation of the components using the measured temperatures and then examine the residual trends to search for evidence of additional physical effects. The resulting errors are believed to be small since the effective étendues a_{fk} are small (typically <1% of the Xcal), their temperatures did not vary much over the course of the flight (∼0.05 K), and there were no systematic differences between the calibration data and the sky data.

The Kirchhoff condition is that $\sum_k a_{fk} = 0$ at each frequency, f. This powerful condition is a statement of detailed balance, or, otherwise expressed, that all of the radiation absorbed by the bolometer must come from somewhere. A way of deriving this condition is to imagine that the entire spectrometer and all its inputs are isothermal. In this case the whole instrument is filled with blackbody radiation, and the motion of the MTM can produce no changes or modulation. This condition can be expressed as $\sum_k a_{fk} P_f(T_k) = 0$, and since all the T_k are the same the result is obtained. We tested this condition by solving for all seven a_{fk} and found that the sum is not significantly different from zero. We therefore impose the condition as a constraint, reducing the number of independent a_{fk} to six.

3.5. Phase

Since the Fourier transform gives complex spectra, in general the a_{fk} are complex. A perfect instrument and properly chosen zero position would make all a_{fk} real. If the zero time in the data is not chosen to be the zero path difference, the a_{fk} will be complex but will all have the same phase, with the phase proportional to frequency. A more complex description of the optics allows a different phase for each of the modeled components. The Xcal and the Ical have nearly the opposite phase, as do the sky horn and the reference horn. The Xcal and the sky horn signals are not in phase with each other, indicating a difference in the spatial distributions of the light produced by the two sources. This is expected because the light from the calibrators must pass through limiting apertures in the horns and is absorbed in the walls, while the emission from the horns originates in the walls and approximately half is emitted following the limiting aperture at the throat of the horn. This difference becomes more important at higher frequencies.

In addition there appears to be a small phase jitter. To correct for this effect each averaged group of interferograms ("coadd") (in both sky and calibration data) has its own small phase adjustment, typically $\psi \sim 10^{-4}$ cm. The largest high frequency corrections are $\sim 0.5°$. For the low frequency data these corrections are entirely negligible but we have retained them to match the calibration data where they are important for the imaginary part of the spectrum.

3.6. Frequency Scale

There are several other optical parameters that must be modeled to achieve photometric consistency and accuracy, including the absolute wavelength (or frequency) scale. The measurements of the mirror position were made with a photo-etched glass scale reference, but the glass scale was calibrated at 295 K and it contracted in cooling to 2 K. In addition, optical beam divergence within the interferometer multiplies the effective path difference by the cosine of the angle of each ray from the principal axis. This effect is substantial because the instrument throughput is large to accommodate diffraction at the long wavelengths under study. The divergence limits the spectral resolution to about 0.8%, but this is overshadowed by the stroke length limitation. These two effects together are calculated to make a 0.5% correction in the frequency scale, which agrees with measurements made of the strong spectrum line from interstellar ionized carbon at 157.7 μmwavelength. This does not match the best fit to a linear photometric model, which is based on observations of the Xcal. The disagreement can be corrected by either a frequency or a Xcal temperature shift. The 0.5% adjustment was chosen for consistency with the measured interstellar spectrum line, with the remainder being absorbed by a 4 mK adjustment in the absolute temperature scale.

3.7. Harmonic Responses

The FIRAS can produce harmonic and subharmonic responses if radiation passes through the instrument more than once, or if a portion of the energy from the beamsplitter is reflected from a fixed component instead of the moving dihedral. Multiple passes are possible because there are substantial reflections at the detectors and filters, and because the beam defining optics are reflective rather than absorbing. Radiation reflected from a bolometer may return to the input concentrating cone and be reflected so that a portion may be received by the detector. This effect is increased by the large aberrations of the off axis parabolic collimators and the concentrating cones, so that $\sim 20\%$ of the radiation from the input fails to reach the output. Moreover, the instrument is not perfectly symmetrical, with an apparent reflectance of the Ical of $\sim 4\%$. If some part of the radiation makes h passes through the interferometer, the phase difference for that part is $2\pi hxf$, where x is the normal path difference for one pass. Therefore the multipass radiation appears in our spectrum at frequency hf. We can turn this around to say that some part of the power we see at frequency f was actually at frequency f/h. Since only a small part of the radiation makes multiple passes, we can approximate each multipass term with the model

$$H_f^h = \epsilon_h Y_{(f/h)} \tag{5}$$

where ϵ_h is the fraction that makes h passes. A second harmonic ($h = 2$) will appear if radiation reflects from the bolometer and goes through the optics again. A half (sub) harmonic ($h = 1/2$) will occur if the radiation on one side reflects from a stationary object rather than the moving mirror. A third harmonic ($h = 3$) indicates 3 trips through the instrument. Each of these terms was tried in a numerical fit, and the second and third harmonics were significant, with a level of $\sim 0.1\%$. This result is consistent with laboratory tests of prototype interferometers and with calculations from the optical design.

3.8. Vibration induced ghosts

The MTM was checked for vibration by timing the sampling pulses derived from the glass reference scale in the mirror mechanism. Two significant resonances were noted at 8.55 Hz and 57.57 Hz. These modes appear to be excited by the MTM velocity reversals and are repeated with about the same amplitude and phase from scan to scan. Vibrations excited by the spacecraft attitude control system wheels were of comparable size but were not coherent with the scan, so they produced no systematic errors. The 57.57 Hz vibration was seen directly in the calibration residuals with an antisymmetric structure centered about the line. A model allowing the position of the mirror to oscillate about the measured position predicts a "ghost" spectrum of just this nature superimposed on the real spectrum. We model this power absorbed at the detector as follows.

There are actually two separate processes which produce sum and difference frequency sidebands. The first frequency modulation occurs when the velocity of the MTM fluctuates, producing sidebands in the interferogram applied to the detector. The second modulation occurs at sampling, when the sample pulses are not equally spaced. These differ because the analog transfer function of the bolometer and amplifier intervene, and the circuit that generates the sample pulses does not track the vibrations exactly. Ideally, the sample pulses would occur at the right moments to cancel the previously generated modulations.

An ideal interferogram of a monochromatic signal is given by

$$I_0(\delta) = I_0 \cos(2\pi\nu\delta), \tag{6}$$

where δ is the path difference measured from the interferogram peak and ν is the frequency in cm^{-1}. Perturbations in the stroke position and sampling time give

$$\delta = vt + \Delta_1 \sin(qt + \alpha), \tag{7}$$

$$t_s = t + (\Delta_2/v)\sin(qt + \beta), \tag{8}$$

where v is the effective scan velocity, t is the time, Δ_1 and Δ_2 are amplitudes of the vibration-induced perturbations, $q = 2\pi f_v$ is the frequency of the vibration, and α and β are phases of the vibrations.

The small signal approximation for the perturbed interferogram is

$$\begin{aligned} I_1(t) &= I_0(\delta) = I_0 \cos(2\pi\nu vt + 2\pi\nu\Delta_1 \sin(qt + \alpha)) \\ &\simeq I_0[\cos(2\pi\nu vt) - 2\pi\Delta_1\nu \sin(2\pi\nu vt)\sin(qt + \alpha)] \\ &\simeq I_0 \left\{\cos(\omega t) + \pi\Delta_1\nu[\cos(\omega t + qt + \alpha) - \cos(\omega t - qt - \alpha)]\right\}, \end{aligned} \tag{9}$$

where $\omega = 2\pi\nu v$ is the angular frequency.

Because the detector and amplifier have different gains and phases at the sum and difference frequencies than at the fundamental frequency, we use (g_+, γ_+), and (g_-, γ_-) for the relative gains and phases. So the sampled output is:

$$\begin{aligned} I_2(t) \simeq\; & I_0 \cos(\omega t) \\ & + I_0\pi\nu\Delta_1[g_+ \cos(\omega t + qt + \alpha + \gamma_+) - g_- \cos(\omega t - qt - \alpha + \gamma_-)]. \end{aligned} \tag{10}$$

When this interferogram is sampled at variable intervals, we obtain additional sum and difference frequency components from the perturbations of the original frequency. In principle there could also be sidebands of the sidebands, but we neglect this effect as second order. The result is

$$\begin{aligned} I_3(t) \simeq\ & I_0(\cos(\omega t) \\ & +I_0\pi\nu\Delta_1[g_+\cos(qt+\omega t+\alpha+\gamma_+) - g_-\cos(qt-\omega t+\alpha-\gamma_-)] \\ & +I_0\pi\nu\Delta_2[\cos(qt+\omega t+\beta) - \cos(qt-\omega t+\beta)]. \end{aligned} \tag{11}$$

The sum and difference terms cancel if $1 = g_+ = g_-$, $0 = \gamma_+ = \gamma_-$, $\Delta_1 = -\Delta_2$, and $\alpha = \beta$. We have compensated the bolometer and amplifier so that $g_+ \simeq g_-$ and $\gamma_+ \simeq \gamma_-$, by using the pole-zero compensator circuit and a Bessel low-pass filter. However, the other two conditions are not met very well because the phase lock loop cannot track the vibrations, and the reference scale does not measure the vibrations at the same locations that affect the optical path difference. Assuming $g_- \simeq 1 \simeq g_+$ and $\gamma_+ \simeq \gamma_- \simeq 0$ gives:

$$\begin{aligned} I_3(t) \simeq\ & I_0\cos(\omega t) \\ & +I_0\pi\nu[\Delta_1\cos\alpha + \Delta_2\cos\beta][\cos(qt+\omega t) - \cos(qt-\omega t)] \\ & -I_0\pi\nu[\Delta_1\sin\alpha + \Delta_2\sin\beta][\sin(qt+\omega t) - \sin(qt-\omega t)]. \end{aligned} \tag{12}$$

In the calibration algorithm, we are most concerned with the even part. By identifying $\Delta = \pi(\Delta_1\cos\alpha + \Delta_2\cos\beta)$, the resulting simplified expression for the correction is:

$$G_f = \Delta(\xi - f)a_{\xi-f,k}P_{\xi-f}(T_k) \tag{13}$$

where ξ is the frequency of the vibration and Δ is the amplitude of the effect. The factor $(\xi - f)$ is the ν above. We allow both negative and positive $(\xi - f)$ to account for the sum and difference terms but use the absolute value in the Planck and étendue subscripts.

Since the blackbody input spectra for the calibration data are known, the amplitude Δ can be fit separately for each calibration interferogram. However, the spectrum of the sky is not known, so this cannot be done for sky data. Instead, a quartic polynomial model of the time variation of Δ was fitted to the calibration data and applied to all of the data. The Δ appeared stable in each ~ 3 day calibration run and the structure of the Δ vs time plot indicates that a fourth order polynomial is sufficient. The 8.55 Hz vibration correction is more stable so a single constant was used to approximate it for the whole mission. These ghosts are more important for the high temperature high frequency data for three reasons: (1) the 57.57 Hz modulation appears at 73 cm^{-1} for slow scans; (2) the factor $(\xi - f)$ is large; and (3) the signal to noise ratio is large enough to make small corrections important. Conversely, the ghost is zero if the input spectrum is nulled, as it nearly is in the low frequency range.

3.9. Detector model

In addition to making an accurate model of the interferometer optics, the operation of the bolometric detectors must be modeled to understand the calibration data. This is particularly true if the calibration is to be valid for the high frequency beyond 30 cm^{-1}, since these regions of the spectrum can be excited only with relatively high temperature calibrators. At high source temperatures direct thermal coupling from the calibrators to the detector housings, combined with increased infrared power incident on the detectors, make major changes in their operating temperatures and hence in their responsivities and time constants.

Each detector is operated with a DC bias voltage applied to a load resistor in series with the detector. The signal itself is measured with a junction field effect transistor (JFET) whose gate is connected to the point between the load resistor and the detector. Since the JFET cannot operate at liquid helium temperature, it is suspended from threads in a small chamber and heated to ~ 70 K by its own bias power.

Nine parameters are used to characterize the bolometer operation. The measured voltage V_m is related to the actual signal voltage V_d and the DC gain and offset of the JFET, J_g and J_0 by

$$V_m = J_0 + J_g V_d. \tag{14}$$

The DC gain is not as close to unity as the AC gain because the JFET offset changes with operating temperature, which is in turn influenced by the DC gate voltage. Next the detector voltage V_d is related to the bias voltage V_b using the resistances of the load resistor, $R_l = 40$ MΩ, and the detector resistance R_d according to

$$V_d = V_b \frac{R_d}{R_l + R_d}. \tag{15}$$

The detector resistance R_d is a function of temperature and the electric field associated with V_d,

$$R_d = R_0 \exp(\sqrt{\frac{T_0}{T_d}})(\rho \frac{V_d}{T_d}) / \sinh(\rho \frac{V_d}{T_d}), \tag{16}$$

where T_d is the detector temperature, and R_0, T_0 and ρ are device parameters. The factor $(\rho V_d/T_d)/\sinh(\rho V_d/T_d)$ is a parameterization of non-linear resistance in the thermometer, and is necessary to achieve a reasonable fit to the DC and AC behavior of the detector. The thermal conductance of the detector to the detector housing $G_{th}(T_d)$ is modeled as a power law,

$$G_{th}(T_d) = G_0 T_d^{\beta_{th}}, \tag{17}$$

and the heat capacity $C(T_d)$ as a polynomial,

$$C(T_d) = C_3 T_d^3 + C_1 T_d, \tag{18}$$

with G_0, β_{th}, C_3 and C_1 as parameters. The steady state heat flow along the leads to the heat sink at temperature T_c is

$$W(T_d, T_c) = \int_{T_c}^{T_d} G_{th}(T) dT, \tag{19}$$

which is used with the bias power and input radiation power to find the DC operating point. From this we follow Mather (1982, 1984a, 1984b) to determine the responsivity, time constant, and noise at any signal frequency. In particular we find that $B(f) = S_0/(1 + 2\pi i f \tau_e)$ where S_0 is the DC responsivity and τ_e is the effective time constant, including all the detector electrothermal feedback effects.

The detectors were mounted in their flight assemblies and tested in a small dewar before installation into the FIRAS instrument. Each assembly included the load resistor, the first stage JFET amplifier, and the concentrator optics. Two sets of tests were performed to obtain the detector characteristics: load curves (voltage as a function of current), and AC electrical transfer functions. The two tests were performed at 12 bath temperatures from 1.5 to 4.2 K. The load curves permit the determination of the all the detector parameters except the device heat capacity, which was determined by measuring the electrical transfer function as a function of frequency. For this test the detector was operated with a small amount of AC voltage signal added to the DC voltage across the load resistor and detector. The transfer function is the ratio of the AC signal appearing across the detector to that applied.

The total of nine parameters with their uncertainties and covariance matrix were determined from these tests. The agreement with the determination of the parameters from the FIRAS in-orbit calibration is poor, with normalized χ^2s of 80 to 800 in various fits for 9 DOF (degrees of freedom). This is probably due to a deficit in the bolometer model.

Since the bolometer is in general nonlinear and the model is linear, a second order effect might be expected. This effect was measured by fitting a second order term to the flight calibration data, with the conclusion that it is not needed. In part this means that the parameters of the optical and bolometer models can be adjusted to partially compensate for this effect. For completeness, we developed an algorithm to make a solution of the nonlinear equations for the bolometer responsivity, based on the time series of measured voltages, but the signal levels seen are too small to require this complexity. We expect the second order terms to be an order of magnitude (signal voltage/bias voltage) smaller than the first order terms, and therefore be very small except when viewing the Moon and the brightest calibration sources.

4. Algorithms

4.1. Data Preparation

In the data analysis we treat the calibration data in the same way as the sky data, with some exceptions to be described. Calibration data are used to derive a calibration model, which is then applied to the sky data and cross checked by direct application to the calibration input data as well.

As the *COBE* orbits the Earth, the FIRAS line of sight traces a path on the sky along a line of approximately constant ecliptic longitude. Each interferogram is a measurement of a certain swath along this line, whose width is equal to the beamwidth. The length of the swath is slightly greater than the beamwidth because of the observing time and the movement of the beam at a rate of 3.5° per minute. The FIRAS beam pattern is designed to be circularly symmetrical and unpolarized, however as the two polarizations proceed on opposite sides of the instrument there could be some polarization dependence. The magnitude of this dependence is likely to be on the same order as the other side to side asymmetries, which is about 1%. The spacecraft rotates at 0.8 rpm throughout the mission, in fact making more than 1/2 revolution while taking a single interferogram. This spin eliminates most of the remaining beam asymmetry and polarization sensitivity. We further reduce the sensitivity to the polarization by averaging data taken in random orientations. The final polarization sensitivity is likely to less than one part in 10000 which would preclude a detection of even a strongly polarized source. To organize the data we divide the sky into 6144 equal area pixels. These are formed by constructing a cube with each face divided into $32 \times 32 = 1024$ squares, projected onto a circumscribed celestial sphere. The projection is not a simple geometrical one, but is adjusted to form equal area pixels having a solid angle of $4\pi/6144$ sr or 6.7 deg^2 (Chan and O'Neill 1975, O'Neill and Laubscher 1976). Because the 7° beamwidth of the sky horn is greater than the separation between pixels (2.6° on average) this binning oversamples the sky and the pixels are not independent.

4.2. Dither subtraction

For each telemetered interferogram, the flight electronics generates a pseudo-random DC voltage that is added to the interferogram before digitization. Its purpose is to exercise a different portion of the analog-to-digital converter range each time, avoiding coherent effects from the nonlinearity of the converter. This DC voltage carries no spectral information and is subtracted immediately in the data processing, which also eliminates the DC component of the spectrum. This DC voltage also adds a transient of known shape at the beginning of the scan as the digital filter "remembers" the previous DC value.

4.3. Coaddition

For both the sky and the calibration data the first step is to assemble a "coadd", i.e. a nominally homogeneous set of data, presumed to be measuring the same thing with the instrument in the same state. Sky coadds are assembled from data taken when the sky horn axis is pointed within a particular pixel at the center of each data collection interval. Calibration coadds are data sets taken when the temperatures of the reference blackbodies and horns are all at stable controlled "plateaus". Note that this results in a fundamental difference between the sky data and the calibration data. A calibration coadd is a continuous (in time) data set spanning up to several hours, whereas sky coadds are data taken in short segments, about 1 minute each, spread out over the ~ 3 day interval that it takes the line of sight to move past each pixel as the spacecraft orbit precesses.

There are several useful results of this coaddition. One is that it reduces the volume of the data from several hundred thousand interferograms to only a few thousand coadds per channel. We also get an estimate of the variance in the data by comparing measurements that are nominally the same. Finally, we use the presumed homogeneity of the ensemble to compare each interferogram to the average, and thus eliminate corrupted or otherwise anomalous data (Isaacman, Read and Barnes 1992).

4.4. Deglitching

Cosmic ray hits on the detectors, or "glitches," are the dominant source of noise for the detectors, but their effects can be reduced by about a factor of 3 to 4 by pulse recognition software on the ground. On average a particle hits a bolometer about 2 times per second, or more than 50 times per interferogram. This rate is close to the unshielded Galactic proton rate of 4 cm^{-2} s^{-1} for our detectors of area 0.5 cm^2. During passages through the

South Atlantic Anomaly the rate can be orders of magnitude higher so all such data are ignored. Each glitch produces an abrupt temperature rise at the detector, and the detector and preamp filters respond with a well defined shape, which for the more energetic (and thus most damaging) hits can be recognized and removed. We remove the glitches using a one-dimensional variant of the radio astronomers' CLEAN procedure (Högbom 1974) to iteratively subtract glitch profiles from each interferogram.

Before executing the CLEAN algorithm, a statistically robust estimate, called a template, of the "correct" interferogram is derived from an ensemble of interferograms by a mid-averaging technique. The mid-average is the mean of the two central quartiles of a distribution, and is calculated separately for each step along the interferogram. This template is itself a good estimator of the interferogram and is immune to the effects of large glitches, but we still improve on its noise level by 10% with the CLEAN procedure. This template is first subtracted from each individual interferogram to get a residual $I_{res}(\delta)$ composed primarily of detector noise and glitches. In some cases the residual interferograms are then seen to have noticeable peaks at the zero path position, indicating real changes of the infrared signal from the mean. For this reason, a second template is created by a statistically robust estimator of $I_{res}(\delta)/I_{res}(0)$.

This second template is then fitted to each residual interferogram with an adjustable amplitude and then subtracted to form a second residual interferogram if the primary and secondary template amplitudes exceed some threshold and the signal to noise ratio for both primary and secondary templates exceeds some threshold. The thresholds are adjusted to make the deglitcher remove glitches at the same rate at the peak of the interferogram as other places in the interferogram.

In the high signal cases the second template is needed to absorb true variations in the input power; without it the deglitcher would subtract real signal at the peak of the interferogram. In near null cases the secondary template can absorb glitches, so real glitches would remain at the peak. By adjusting the thresholds the optimum threshold is achieved. It is recognized when the apparent glitch rate is the same at the peak as elsewhere.

A modified CLEAN algorithm is then executed on this second residual until all glitches larger than 3.5 σ have been removed. In this context σ is defined as 1.25 times the median absolute value of the residual noise. Since cosmic ray glitches are observed to have an E_g^{-3} differential number distribution, where E_g is the glitch amplitude, residual events are numerous but small and there is little to be gained by more sophisticated algorithms. A typical interferogram has 10 recognizable glitches removed, about 2% of the total number of points in each interferogram.

Next, all interferograms are compared with the ensemble midaverage, and any that still differ from the midaverage by over 50% in overall noise level or have more than 6 points beyond 6 sigma are eliminated. This process, called consistency checking, removes data that have been contaminated by crosstalk from the spacecraft electronics, variations in the calibrator temperatures, errors in the commands, and so forth. Approximately 1% of the data are lost this way. The remaining interferograms are averaged together, and the primary and average of the secondary templates are added back to form an "average" interferogram or "coadd". The individual interferograms are also Fourier transformed into spectra to determine the rms variation for each frequency. The rms values represent the spread in the data and are used to estimate the error of the mean spectrum. A study of the noise correlations between frequencies showed that 99% of all the correlations could be explained by the apodization, and thus the non-Gaussian nature of the residuals of the glitches is not important.

When the interferograms are averaged together, the engineering variables that describe them are also averaged together, so that a mean bias voltage, temperature, and attitude set are derived for each coadd.

The CLEAN algorithm was tested by varying its thresholds and by comparing its outputs with the first templates, which are quite different statistically robust estimators of the best interferograms. There was no systematic difference found, but the residual noise level of the CLEAN algorithm was better than that of the templates, as determined from consistency of separate coadds. The same deglitching algorithm and control parameters are used for the calibration data and the sky data, so no systematic error should result. The largest effect was a difference of about 30μK on the fitted absolute temperature adjustment, but this is negligible compared to the overall shift of -7.4 mK.

4.5. Baseline curvature subtraction

The interferograms for high signal levels are observed to be superposed on a curved baseline. This baseline curvature is produced by a change of the focus of the instrument as the path difference is changed, which then modulates the background power on the detector. To extend the accuracy of the instrument to the lowest possible optical frequencies, the baseline curvature must be removed in a physically realistic manner. Measurements showed that the baseline curvature is well represented by a quartic polynomial, so we fit and subtract such a polynomial from each coadd. Although this is a simple form of digital high pass filtering, it has wider importance. The low frequency modulation represented by the baseline curvature produces responses at higher frequencies when the baseline is apodized and Fourier transformed, and is effectively a zero frequency spectrum ghost with sidelobes. The digital transient is also fit and subtracted at this time. The final "average" interferogram is then apodized and Fourier transformed. The result is a single complex spectrum for each channel, scan mode, and pixel with an associated variance spectrum.

4.6. Summary of calibration procedure

To review, the input is an inverse Fourier transform of the apodized coadded interferogram. A model of the electronics, Z_f, is used to remove the filter function of the electronics. Next, the state of the bolometer is found from the temperatures and other engineering data. Then a linear model of the bolometer B_f is used to infer the radiation spectrum, R_f, on the bolometer. Since these are done with a discrete spectrum we use an index to denote the frequency, $R_f = Y_f/(Z_f B_f)$, where we use the notation Y_f for the discrete Fourier transform of the input interferogram. The harmonics, vibration corrections and emission from all of the known components, $a_{kf}P_f(T_k)$, in the beam are subtracted from the data. The result is divided by the étendue derived from the Xcal, as described in (2) above. Finally a small phase adjustment, ψ, is made to minimize the imaginary part of the signal. The phase adjustment is only significant for the high frequency data when observing the Galaxy. Thus,

$$S_f = e^{if\psi}(R_f - \sum_k a_{fk}P_f(T_k) - H_f - G_f)/a_{fx}. \tag{20}$$

The net result, S_f, is presumed to be the sky spectrum.

5. Calibration parameters

5.1. Available calibration data

The detector and JFET parameters, the constants of the filters, and the complex étendues of the components are needed to calibrate the spectra. Some of these are obtained from data taken before the launch, where both the input and the outputs could be measured, but most are derived from in-orbit data. The Xcal was commanded into the horn 16 times during the flight, typically for 3 days, and data were taken while the temperatures of the controllables were set at various temperature combinations. For data taken at these times there is no "unknown" radiation coming down the sky horn, only the blackbody radiation from the Xcal, which is at a known temperature. Under these conditions the remaining spectra, after all known sources are removed, should be null. All of these parameters are fit by minimizing a χ^2, which is the weighted sum of the square of all of the residuals and constraints.

5.2. Temperatures

We describe here the criteria for selecting the temperature combinations used in the calibration, and then the method for adjusting the temperatures to fit the data. High temperatures give large signals from which it is easy to determine the instrument gains, while temperatures near 2.7 K approximate the sky data. Both are necessary for accuracy.

For the low frequency data, calibration temperatures above ~ 6 K are not needed and were not used. Approximately 82% of the calibration data were taken with all of the controllables near 2.7 K, reproducing most sky observing conditions. For the high frequency data (>20 cm^{-1}) this approach could not be used, since there is insufficient radiation at frequencies above $\sim 5\,\mathrm{cm}^{-1} \times (T/1\,\mathrm{K})$, where T is the Xcal temperature. For these frequencies, temperatures up to 25 K were used. For such high temperature sources, the bolometer operating point is significantly different from that used for sky observations, so we depend on the accuracy of the bolometer model and the ability of the calibration data to verify that accuracy. The result is that at the higher frequencies the gain of the instrument for the calibration may not match that for the sky. The zero point is well determined for all frequencies because most of the calibration data were taken near the 2.7 K null condition. To further reduce the effect of the high temperatures we have subtracted the average of the residual of all of the 2.7 K null calibration data from the sky data.

The photometric calibration was better than the temperature measurements. Digitization and other noise limited the measurement precision to ~ 0.2 mK at 2.7 K, with much larger errors at higher temperatures. Because the precision of the radiometer is better than this, these temperature errors increased the χ^2 beyond what should be expected. Therefore the temperatures of the controllables in each calibration data set themselves were treated as parameters, dT_{kn} to be fit. This did not change the number of degrees of freedom since each temperature measurement was still treated as a valid measurement. The difference between the measurement and the fit was added into the χ^2 along with the other errors, using the form $\sum_{nk} dT_{kn}^2/\sigma_T^2$. For this purpose we assumed that the statistical uncertainty of a temperature measurement was $\sigma_T = 0.26\,\mathrm{mK}\,(T/2.7\,\mathrm{K})^3$, or approximately the resolution of the thermometer analog to digital converter. This is not to be confused with the accuracy of the thermometers, which is not important except for the Xcal.

Through this procedure we found a systematic trend in the fitted Ical temperature. The magnitude of this trend is ~ 6 mK with most of it occurring in the first 60 days of the mission. The trend was observed both in the calibration data and in the sky data. Therefore an exponential drift of the Ical thermometer was included in the fit to remove this trend which was seen in all of the channels. The time constants found were 60-100 days depending on the channel. The source of this trend was not determined although it was clearly in the instrument since it was observed both with the Xcal in and out of the horn. The sky data were used to remove it since those data have a larger signal to noise ratio than the smaller calibration data set.

An important issue is the absolute accuracy of the thermometers on the Xcal. As a test, a fit was done with an arbitrary offset in the Xcal thermometer and the result was -7.4 ± 0.2 mK for the offset. This is considerably larger than the ~ 1 mK expected from the pre-flight calibration of the thermometers. Either some unknown systematic effect is perturbing the fit or the thermometers are not calibrated as well as was initially thought. To cover both possibilities we have set the offset to a middle value of -4 mK and assign a 5 mK $1\,\sigma$ error to the accuracy of the Xcal thermometers.

5.3. Bolometer and other nonlinear parameters

We allow the adjustment of a variety of parameters, θ_j, which enter the fit in a nonlinear fashion. These include the R_0, G_0, ρ, C_1, J_0, and J_g of the detector model previously discussed. The bolometer parameters were measured before launch, but the other parameters were not. When *a priori* constraints exist, these are included in the χ^2 in the form of the term $\sum_{jm}(\theta_j - \theta_{j0})W_{jm}(\theta_m - \theta_{m0})$. The W_{jm} are taken from preflight measurement error estimates. Other unconstrained parameters are also included in the fits: the J_0 just mentioned; the ϵ_h and Δ from the harmonic and vibration models; and the Ical offset temperature B_I. These are all included in the vector θ, but for the unconstrained ones (those without preflight data), the W_{jm} are zero.

Our approach to finding the étendues, phases, and bolometer constants is nothing more than a least squares fit to all the calibration data. It is not a trivial computation since the number of parameters is large $(2500 - 4000)$ and some of them do not enter linearly. Nevertheless, the random errors are known (estimated from the deviations within the coadds), and the χ^2 can be calculated in a straightforward manner as a sum over frequencies f, observations n, and emitting components k:

$$\chi^2 = \sum_{fn} \left| \frac{Y_{fn} e^{if\psi_n}}{Z_f B_{fn}} - \sum_k a_{fk} P_f(T_{kn} + dT_{kn}) - H_{fn} - G_{fn} \right|^2 / \sigma_{fn}^2$$

$$+\sum_{fk}\frac{a_{fk}^2}{\sigma_{a_{fk}}^2}+\sum_{n}\frac{\psi_n^2}{\sigma_\psi^2}+\sum_{nk}\frac{dT_{kn}^2}{\sigma_{T_{kn}}^2}+\sum_{jm}(\theta_j-\theta_{j0})W_{jm}(\theta_m-\theta_{m0}). \tag{21}$$

Temperature adjustments dT_{kn} are typically 0.1 mK for near 2.7 K temperatures. Harmonic corrections described by equation (5) above are included in H_{fn}, and vibration corrections described by equation (13) are included in G_{fn}. Since some of the components (bolometers, dihedral mirrors) do not get hot enough to fit the high frequency emission there are the added constraints on the étendues scaled by $\sigma_{a_{fk}}$. These were set at 1% of the Xcal a_{fk} for the dihedral and the collimator and 10% for the bolometer. The other étendues were unconstrained. The constraint on ψ, σ_ψ, is set at 5×10^{-4} cm to allow phase adjustment only when there is sufficient signal to noise to warrant a correction.

5.4. Fitting procedure

The process of fitting the instrument model to the calibration data is complicated by the large number of parameters. Conventional fitting algorithms require the construction and inversion of the Hessian of χ^2 with respect to the parameters. There are also algorithms that do not use the Hessian, but those that do have two advantages: 1) the Hessian provides scale information that is not present in the gradient, which enables such algorithms to be more robust and make more progress per iteration; and 2) the inverse of the Hessian is the covariance matrix of the parameters, which is necessary for error analysis. However, the inversion of a matrix of such large rank ($\simeq$ 4000) is not generally tractable. We were able to invert this matrix by taking advantage of its special form; specifically, that the elements of the Hessian corresponding to two temperature correction parameters from different coadds are zero. This made the inversion possible, though still not speedy, so the method we adopted was to use a non-Hessian method (conjugate gradient) until we thought we were in the vicinity of the solution, and then proceed with a Hessian method (Levenberg-Marquardt) to refine the solution and produce the covariance matrix (Press *et al.* 1986).

6. Error Discussion

The normalized χ^2 resulting from this fit is 2.8218 (27873 DOF) for the left low detector, short slow stroke data ($2.27<\nu<21.54$ cm^{-1}), and 4.53 for (159353 DOF) for the right high detector, short slow stroke data ($2.27<\nu<96.28$ cm^{-1}). We concentrate on these two data sets since they contain 85% of the weight of the total data and they are typical of the other data; the χ^2 and DOF for each channel/scan mode combination are summarized in Table 2. Although the χ^2 is larger than statistically expected, the errors are concentrated both in the high temperature data (which is not representative of the sky data), and in the imaginary components (which contain only the fine spectral resolution in the sky data). If only the 2.7 K null data are considered the χ^2/DOF is 1.15 for the high frequency data and 1.37 for the low frequency data. Furthermore, the calculated relative étendues a_{fk}/a_{fx} agree to better than 0.1% for the two low frequency detectors and 0.5% for the high frequency detectors. Even the detector parameters, which are fit separately for the short scan and the long scan data, agree to better than 2%. This agreement, along with the reasonable χ^2 for the sky-like null data, gives us confidence in the models. A typical residual spectrum is shown in Figure 8.

Further tests of the calibration are obtained by searching the calibrated map of the sky for features relating to changes of the instrument state. The largest such changes occurred during the time from May to August 1990. In this time period, it was impossible to keep both the Earth and the Sun below the Sun screen, and the Earth illuminated the top of the instrument during part of the orbit. The data taken with the Earth above the instrument were rejected from the maps, but the thermal transient produced by the heat from the Earth was large and long. As a result, we raised the set point of the horn temperature controllers as high as 6 K to achieve stability. The calibration of the horn étendues is sufficiently accurate that there is no recognizable trace of this event in the sky maps, at the level of $\sim$ 0.1 mK in the CMBR temperature.

7. Error Analysis

The calibrated spectra $S_{fn} = S(Y_{fn}, \theta, a_{fk}, T_{kn})$ are functions of the calibration parameters, and their uncertainties can be determined from their partial derivatives with respect to those parameters, combined with the covariance matrix for the fitted parameters. Any systematic error which has been modeled and included in the fit is included in this general form. Those for which no model has been fitted must be treated separately. The formula is as follows:

$$\delta S_{fn} = \frac{\partial S}{\partial Y}\delta Y + \frac{\partial S}{\partial a_{fk}}\delta a_{fk} + \frac{\partial S}{\partial \theta}\delta\theta + \frac{\partial S}{\partial T_{kn}}\delta T_{kn} + (\frac{\partial S}{\partial a_{fk}}\frac{\partial a_{fk}}{\partial T_x} + \frac{\partial S}{\partial \theta}\frac{\partial \theta}{\partial T_x})\delta T_x \tag{22}$$

The covariance of the calibrated sky spectra S_f has contributions from four statistically independent sources: detector noise in the uncalibrated spectra Y_{fn}, the covariance of the fitted calibration parameters θ, a_{fk}, random errors in the measurement of the temperatures of the controllables T_{kn}, and the absolute calibration error in the Xcal thermometry T_x.

7.1. Detector noise

The uncertainty in the uncalibrated spectra δY_f is estimated from the scatter within each coadd. This variance is principally particle hits and detector noise. The covariance between different frequencies due to apodization of the interferogram prior to Fourier transformation, is ∼5%, and we neglect it. Each calibrated spectrum S_f is a function of one uncalibrated spectrum and the fitted calibration parameters. This dependence is described by (20). The propagation of the error in Y_f to the covariance in S_f is straightforward and will be described in the next sections. In the following the covariance in S_f due to detector noise is called D. This covariance is zero between different pixels. There is a small covariance between different frequencies due to the vibration and harmonic corrections in addition to the apodization, but these are negligible. Not negligible are residual systematic effects at the vibration frequencies. To account for these the detector noise was multiplied by ten for the bins contaminated by vibration.

7.2. Fitted calibration parameter covariance

The parameters that are determined by fitting to the calibration data can be divided into three groups: miscellaneous parameters θ, including bolometer parameters and others, of which we fit 14; the complex parameters a_{fk} associated with each frequency, of which there are between 400 and 2000 depending on channel and scan mode; and the parameters d associated with each calibration coadd, which include four temperature corrections and a phase correction, for a total of between 750 and 2400 parameters, again depending on channel and scan mode. The parameters θ and a_{fk} are used to calibrate the sky data, but the d parameters are not. This is because each d parameter pertains to a particular calibration coadd, and therefore provides no information about the sky coadds. Therefore only that portion of the covariance matrix from the fit which pertains to the θ and a_{fk} parameters affects the covariance of S_f.

Thus it is convenient to partition the the covariance matrix of the parameters along these three sets. The calibration program computes the Hessian of χ^2 with respect to the parameters, which can be written as the block matrix

$$\begin{pmatrix} P & Q & R \\ Q^T & E & F \\ R^T & F^T & D \end{pmatrix}$$

where P is the $\theta\theta$ part of the matrix, Q is the θa_{fk} part, R is the θd part, etc. [One fact about this matrix that is important later is that the matrix E (the $a_{fk}a_{f'k'}$ part of the Hessian) is block diagonal on frequencies, or in other words it does not mix frequencies. One consequence is that E, although quite large, is easily inverted.] The inverse of the full Hessian matrix is the correlation matrix of the fitted parameters,

$$\begin{pmatrix} C_P & C_Q & C_R \\ C_Q^T & C_E & C_F \\ C_R^T & C_F^T & C_D \end{pmatrix},$$

which has been partitioned in the same way. Therefore the part of the correlation matrix of the calibrated sky spectra S_f due to the correlation of the calibration parameters is

$$\mathrm{cov}(S_f S_{f'}) = \frac{\partial S_f}{\partial \theta} C_P \frac{\partial S_{f'}}{\partial \theta} + \frac{\partial S_f}{\partial \theta} C_Q \frac{\partial S_{f'}}{\partial a} + \mathrm{sym} + \frac{\partial S_f}{\partial a} C_E \frac{\partial S_{f'}}{\partial a}, \tag{23}$$

where it should be remembered that each S carries its own index of sky position, which is suppressed here; therefore the S_f and $S_{f'}$ are in fact not interchangeable even when $f = f'$. The notation "sym" means symmetrizing the previous term with respect to the two S. This equation gives us $\mathrm{cov}(SS)$ exactly, but the matrices involved are quite large, which would make the calculation of a particular $\mathrm{cov}(SS)$ rather time-consuming. Therefore a more convenient approximation is desirable.

To find a more convenient form, we will use the identities

$$C_E = E^{-1} + E^{-1}\left[Q^T C_P Q + Q^T C_R F^T + F C_R^T Q + F C_D F^T\right] E^{-1} \tag{24}$$

and

$$C_Q = -\left[C_P Q + C_R F^T\right] E^{-1}. \tag{25}$$

These identities result from the matrix equivalent of inversion by cofactors; they are straightforward (albeit quite tedious) to verify. Substituting these into (23), we get

$$\begin{aligned} \mathrm{cov}(S_f S_{f'}) = {} & \frac{\partial S_f}{\partial \theta} C_P \frac{\partial S_{f'}}{\partial \theta} - \frac{\partial S_f}{\partial \theta} C_P Q E^{-1} \frac{\partial S_{f'}}{\partial a} - \mathrm{sym} \\ & + \frac{\partial S_f}{\partial a} E^{-1} Q^T C_P Q E^{-1} \frac{\partial S_{f'}}{\partial a} + \frac{\partial S_f}{\partial a} E^{-1} \frac{\partial S_{f'}}{\partial a} \\ & - \frac{\partial S_f}{\partial \theta} C_R F^T E^{-1} \frac{\partial S_{f'}}{\partial a} - \mathrm{sym} \\ & + \frac{\partial S_f}{\partial a} E^{-1}\left[Q^T C_R F^T + F C_R^T Q + F C_D F^T\right] E^{-1} \frac{\partial S_{f'}}{\partial a}. \end{aligned} \tag{26}$$

We omit the terms in the last two lines of this equation. We have indications that these terms are negligible, though this is undergoing further study.

Now define Λ to be the matrix square root of C_P

$$\Lambda^T \Lambda = C_P, \tag{27}$$

which exists because C_P is a positive definite symmetric matrix. Define the quantity

$$J_f = \Lambda \frac{\partial S_f}{\partial \theta} - \Lambda Q E^{-1} \frac{\partial S_f}{\partial a}. \tag{28}$$

Then the correlation of the sky spectra is

$$\mathrm{cov}(S_f S_{f'}) \approx \frac{\partial S_f}{\partial a} E^{-1} \frac{\partial S_{f'}}{\partial a} + J_f^T J_{f'} \tag{29}$$

where the approximation refers to the omission of the last two lines of (26); this is an exact expression for the first two lines. In the following we will discuss each of these terms in turn.

The first term in (29), which we call the PEP term, is not diagonal in pixels, but it is diagonal in frequency, because the a_{fk} are used for all pixels, but calculated separately (approximately) for each frequency. Thus E is diagonal in frequency, and S_f only depends on $a_{f'k}$ if $f' = f$. Further,

$$\frac{\partial S_f}{\partial a_{fk}} = \frac{P_f(T_k) - P_f(T_{\mathrm{bolo}})}{a_{fx}} \tag{30}$$

for k not the Xcal, where $P_f(T)$ is a Planck I_ν, T_k is the temperature of the kth emitter, and T_{bolo} is the bolometer temperature; in the case k is the Xcal, $P_f(T_k)$ must be replaced by S_f. Thus this term, except for the a_{fx} derivative, depends only on the temperatures of the various emitters in the instrument. Almost all the sky data are quite near the same temperature conditions. These are: Ical 2.76 ± 0.006 K, MTM 2.0 ± 0.4 K, horns 2.75 ± 0.005 K, mirrors 1.56 ± 0.02 K, and bolometers 1.52 ± 0.017 K. Furthermore, in the low frequency channel outside the Galaxy, S_f is well approximated by a Planck spectrum $P_f(T_{\mathrm{CMBR}})$ independent of pixel, so for this data PEP is approximately independent of the pair of pixels at which the two S_f are located.

The second term in (29), which we name JCJ, includes errors that affect the whole spectrum and all pixels. Thus it is not diagonal in frequency, nor is it diagonal in pixels. At a particular pixel the matrix J has frequency indices on one side, and on the other side a number of indices equal to the number of p parameters, of which there are 14 (though the index actually refers not to p but to the eigenvectors of C_P). Thus it can be viewed (at each pixel) as a relatively small number of spectra. The dependence on pixel comes about as a dependence on the instrument state and the sky spectrum. If we use a typical instrument state, and look at a set of typical pixels (Galactic pole, Galactic center, etc.) we can examine the behavior of J under various observing conditions. We find in each case that there is one (or sometimes two) dominant spectra in J, that is, one J spectrum has a square norm that is most of the sum of square norms of all the J spectra. In this case, we can simply use the outer product of this J spectrum with itself to adequately approximate the $J^T J$ term.

7.3. Random Temperature Measurement Errors

Although the temperature correction parameters derived from the calibration fit cannot be used to correct the sky spectra, they do give us information about the covariance U of random deviations of the temperatures of the emitting components from their measured values. The calibration fit shows no correlation between temperature corrections to different coadds, so we conclude that the temperature deviations are not correlated between pixels. The result of propagating this covariance to the sky spectra is called the PUP term. From (20) we see that

$$PUP = \frac{1}{a_{fx}a_{f'x}} \sum_{k,k' \neq \mathrm{Xcal}} a_{fk} \frac{\partial P_f(T_k)}{\partial T_k} U_{kk'} a_{f'k'} \frac{\partial P_{f'}(T_{k'})}{\partial T_{k'}}. \tag{31}$$

This term depends only on the temperature state of the instrument; in particular, it is independent of the sky spectrum. As noted in the discussion of PEP, most sky observations were done in essentially the same state, so PUP is largely independent of pixel. It can be further simplified by noting that the dominant term in most cases is the Ical temperature error, since the Ical has the largest étendue. In fact a fair approximation to PUP is

$$PUP \approx \frac{\partial P_f(T_{\mathrm{Ical}})}{\partial T_{\mathrm{Ical}}} \sigma^2_{\mathrm{Ical}} \frac{\partial P_{f'}(T_{\mathrm{Ical}})}{\partial T_{\mathrm{Ical}}}, \tag{32}$$

since $a_{f\mathrm{Ical}} \sim -a_{fx}$. Note that the errors here are correlated between frequencies, but not between pixels.

7.4. Absolute calibration error in Xcal thermometry

There is an error due to the uncertainty in the absolute temperature scale, independent of the random errors discussed in the preceding subsection. We call this term PTP. This produces a term similar to PUP, but with only one temperature variance:

$$PTP = \frac{\partial P_f(T)}{\partial T}(\delta T_x)^2 \frac{\partial P_{f'}(T)}{\partial T} \tag{33}$$

where δT_x is the uncertainty in the temperature of the Xcal at the relevant temperature. We are most concerned with the temperature of the CMBR, and therefore with the thermometer error at that temperature. As discussed in the section on the thermal model, the internal consistency of the 3 Xcal thermometers was approximately 3 mK at 2.7 K, but the experiments to confirm the absolute temperature scale have not been completed. As noted in the discussion of the frequency scale adjustment, we found a discrepancy of $\sim 0.3\%$ between the photometrically derived frequency scale and the observations of the interstellar spectrum line of ionized carbon. If this were due to an error in the temperature scale it would decrease the 2.7 K measurement by ~7.4 mK. Such an error would not affect our ability to recognize a distortion from the Planck form, but only our measurement of the absolute temperature. For our final estimate we noted that the discrepancy between the thermometers and our calibration model was −7.4 mK. Since this is larger than the estimated statistical errors of either we have chosen a midpoint adjustment of −4 mK to the Xcal thermometry and estimate the error as 5 mK. To get a 95% confidence limit, we take twice this error, but it is based on a single degree of freedom and is not itself precise.

7.5. Summing the errors

These errors add in quadrature, which is to say the total covariance is just the sum of all of the covariance parts discussed in the previous sections:

$$\mathrm{cov}(S_f S_{f'}) = D + PEP + JCJ + PUP + PTP. \tag{34}$$

In principle this sum is easy to do. In practice however the resulting matrix would contain $\sim 10^{12}$ elements, none of which are zero, and would be neither convenient nor illuminating. Instead we keep the terms separate (with $\sim$ 1000 numbers). Figure 9 illustrates typical values of each term. For most interesting calculations only one or two of the terms dominate, and the rest can be ignored. Which terms dominate depends on the application.

7.6. Other errors

The FIRAS is a differential instrument, calibrated in situ with a full beam blackbody calibrator at the entrance aperture. The calibrator is at a known temperature (subject to the uncertainties discussed above), it is isothermal (within 1 mK), it is black (the reflectance is of order 3×10^{-5}), and it covers the entire aperture (the leakage around the edge is less than 10^{-4} and the sky is of similar temperature). As a result, direct substitution of the calibrator for the sky reduces the effects of almost every error of instrument design, calibration procedure, or software. This reduction is particularly effective when the sky spectrum resembles a blackbody at the calibrator temperature, as it does for frequencies below $20\mathrm{cm}^{-1}$.

Nevertheless, we were concerned over other possible sources of error in the FIRAS calibration procedure. There are several tests which might uncover errors of certain types. Software errors were minimized by verifying each piece of software by someone other than the software author, and validating it by checking both input and output for selected cases. All bugs uncovered were corrected. Also where practical, the output of programs written in different languages were compared with the same input.

Algorithms were checked by systematically altering the input (by adding a known distortion) and verifying that it appear on the output. These checks confirmed the software to the bit errors in the computer or about 10 digits (more than the least significant bits are lost because there are many computations) which is 5 digits more than are required for the accuracies reported.

When sky maps were prepared there were several tests done. First, in each detector scan mode combination parts of the sky were viewed several times. The variation in these different observations was compared with the prediction from the detector noise. In all detector scan mode combinations the normalized χ^2 was ~ 1.2 in agreement with the χ^2 in the calibration solutions. This higher estimate of noise was used on the final product. It is now known that a factor of 1.16 is due to the apodization, so that the unexplained portion of the normalized χ^2 is only a factor of 1.04. Second, maps from different detector scan modes were compared, and the differences were consistent with those expected from noise alone. This very powerful test shows that completely separate runs of the calibration software, deglitcher, and all other parts of the analysis yield the same answer.

Three known problems remain unresolved. First, the detector model may be inadequate at the high emitted power and high detector temperature of the high frequency calibration. This is suspected from the behavior of the phase shifts of the high temperature calibrations. To account for this, the gain error estimate produced from the internal consistency of the solution has been inflated by a factor of 10. Second, as previously described, the photometric model and the thermometers disagree by 7.4 mK. To account for this, we weight them approximately equally and state a systematic temperature uncertainty of 10 mK. Thirdly, the effects of systematic vibrations are not entirely removed by the model. To account for this, we have increased the error bars at frequencies affected by vibrations by a factor of 10.

8. Conclusion

From the internal consistency of the calibration, we find that the calibration data are good to less than 0.1% of the peak of the CMBR for relative measurements. The dominant systematic error in the high frequency channels is a gain uncertainty, which is less than about 3%. The limits on most individual pixels are set by the noise of the detector. The absolute temperature accuracy is limited by the error on the Xcal thermometry. The high frequency gain error is limited by the calibration error.

Although the calibration model could be improved (particularly for high signal levels), the fact that χ^2/DOF is near 1 for the sky data implies that it is unlikely that there would be a significant improvement in the final errors in the calibrated data.

9. Acknowledgements

The calibration of the FIRAS was an exceptionally challenging software and data analysis task. The massive data sets were prepared and examined very carefully by a team which included S. Brodd, J. Gales, N. Gonzales, S. Macwan, L. Rosen, F. Shuman, and H. Wang. The NASA Centers for Computational Sciences generously donated time on their Convex C3240. The engineering staff at NASA GSFC which built the instrument so well was led by R. Maichle and M. Roberto, and M. Ryschkewitsch, J. Webb, and J. Wolfgang provided critically important leadership at times of crisis in the launch preparations. This instrument could not have succeeded without their combined efforts.

	Low Frequency		High Frequency	
	short	long	short	long
Samples(fast)	2	8	2	2
Samples(slow)	3	12	3	3
Frequency Interval	0.57	0.14	0.57	0.57
Frequency Range	2.8-20	0.7-20	2.8-96	2.8-96

Table 1: Number of raw samples averaged by on-board processing in each scan mode/frequency combination. Frequencies are in cm^{-1}

	Low Frequency		High Frequency	
	Left	Right	Left	Right
Short Slow	78653	58102	616665	722597
DOF	27873	27941	165179	159353
Long Fast	50728	48899	202482	291555
DOF	37689	37689	75917	74361

Table 2: χ^2 and number of degrees of freedom for each channel/scan mode combination.

REFERENCES

Bennett, C. L., & Hinshaw, G. 1993, in AIP Conf Proc 278, *Back to the Galaxy*, ed. S. S. Holt & F. Verter, (New York: AIP), 257.

Boggess, N. W. *et al.* 1992, ApJ, 397, 420

Chan, F. K., & O'Neill, E. M. 1975, Feasibility Study of a Quadrilateralized Spherical Cube Earth Data Base, (EPRF Technical Report 2-75) (NTIS Report AD-A010232) (Computer Sciences Corp. CSC/TR-75/6007)

Fixsen, D. J. *et al.* 1994, ApJ, 420, 445

Hemmati, H., Mather, J. C., & Eichhorn, W. L. 1985, Appl. Opt., 24, 4489

Högbom, 1974, A&AS, 15, 417

Isaacman, R., Read, S., & Barnes, W. 1992, in *Astronomical Data Analysis Software and Systems I* ed. D. M. Worall, C. Biemesderfer, & J. Barnes, (Astronomical Society of the Pacific Conf. Ser., Vol. 25) p. 403

Mather, J. C. 1982, Appl. Opt., 21, 1125

Mather, J. C. 1984a, Appl. Opt., 23, 584

Mather, J. C. 1984b, Appl. Opt., 23, 3181

Mather, J. C., Toral, M., & Hemmati, H. 1986, Appl. Opt., 25, 2826

Mather, J. C. *et al.* 1990, ApJ, 354, L37

Mather, J. C. *et al.* 1994, ApJ, 420, 439

O'Neill, E. M. & Laubscher, R. E. 1976, Extended Studies of a Quadrilateralized Spherical Cube Earth Data Base, (NEPRF Technical Report 3-76) (NTIS Report AD-A026294) (Computer Sciences Corp. CSC/TR-76/6008)

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1986, Numerical Recipes (Cambridge: Cambridge University)

Wright, E. L. *et al.* 1991, ApJ, 381, 200

This preprint was prepared with the AAS LATEX macros v4.0.

Figure 1. Simplified optical layout of FIRAS, showing the positions of the horns, Xcal, Ical, mirrors, grids, MTM and detectors. The drawing is not to scale, the sky horn is 2 m long and the other parts fit in a box $\sim$ 50 cm on a side.

Figure 2. Typical interferogram, with the Ical $\sim$ 10 mK warmer than needed to match the sky. The peak at position $\sim$ 355 can be nulled within the detector noise level. FIRAS produced $\sim$2 million such interferograms.

Figure 3. Cutaway drawing of the Xcal. The diameter of the Xcal is $\sim$ 20 cm. Light striking the cone suffers 7 specular reflections before escaping from the cone.

Figure 4. Antenna pattern for the FIRAS horn as measured on the ground before launch at 3 cm^{-1}.

Figure 5. Transfer function of the post detector electronics for the left low frequency channel: a) gain, b) phase.

Figure 6. a) Étendue from Xcal to left low detector, in short slow scan mode. The ideal interferometer would have a value of 0.75 cm^2 sr. b) Ratio of Ical étendue to Xcal. c) Ratio of sky horn (solid line) and reference horn (dashed line) étendues to Xcal. d) Ratio of other étendues to Xcal: dihedral mirrors (solid line), structure (dashed line), and bolometers (dotted line).

Figure 7. Apodization function used for the short strokes. a) apodization function, b) apodization function folded at $\delta = 0$ and summed, c) spectral window function, real part (solid line) and imaginary part (dashed line).

Figure 8. Typical residual after the calibration model has been subtracted. The spectrum here is from the average of 93 interferograms to reduce the noise. For comparison, a 2.7 K blackbody has a peak brightness of 120×10^{-6} erg cm^{-2} s^{-1} sr^{-1} cm.

Figure 9. Various error terms for a typical point in the sky. See the text for the definition of each error term. a) Low frequency, b) high frequency.